TECHNISCHE UNIVERSITÄT DRESDEN

Center for Information Services and High Performance Computing (ZIH)

Performance Analysis for GPU Accelerated Applications

Working Together for more Insight

Willersbau, Room A218

Tel. +49 351 - 463 - 39871

Guido Juckeland (guido.juckeland@tu-dresden.de)

Outline

- Motivation
- Performance Analysis 101
- Using Performance Tools for Accelerators
- Examples
- Summary & Outlook

MOTIVATION

Many High-Noon Situations

I know, what my code does!

Use my system efficiently!

User

System Provider

Performance tools can provide an objective view

Many High-Noon Situations (2)

I need more information!

Why do you care?

Tool Developter

Hardware Vendor

Reaching Higher with Cooperation

PERFORMANCE ANALYSIS 101

What do you want to know?

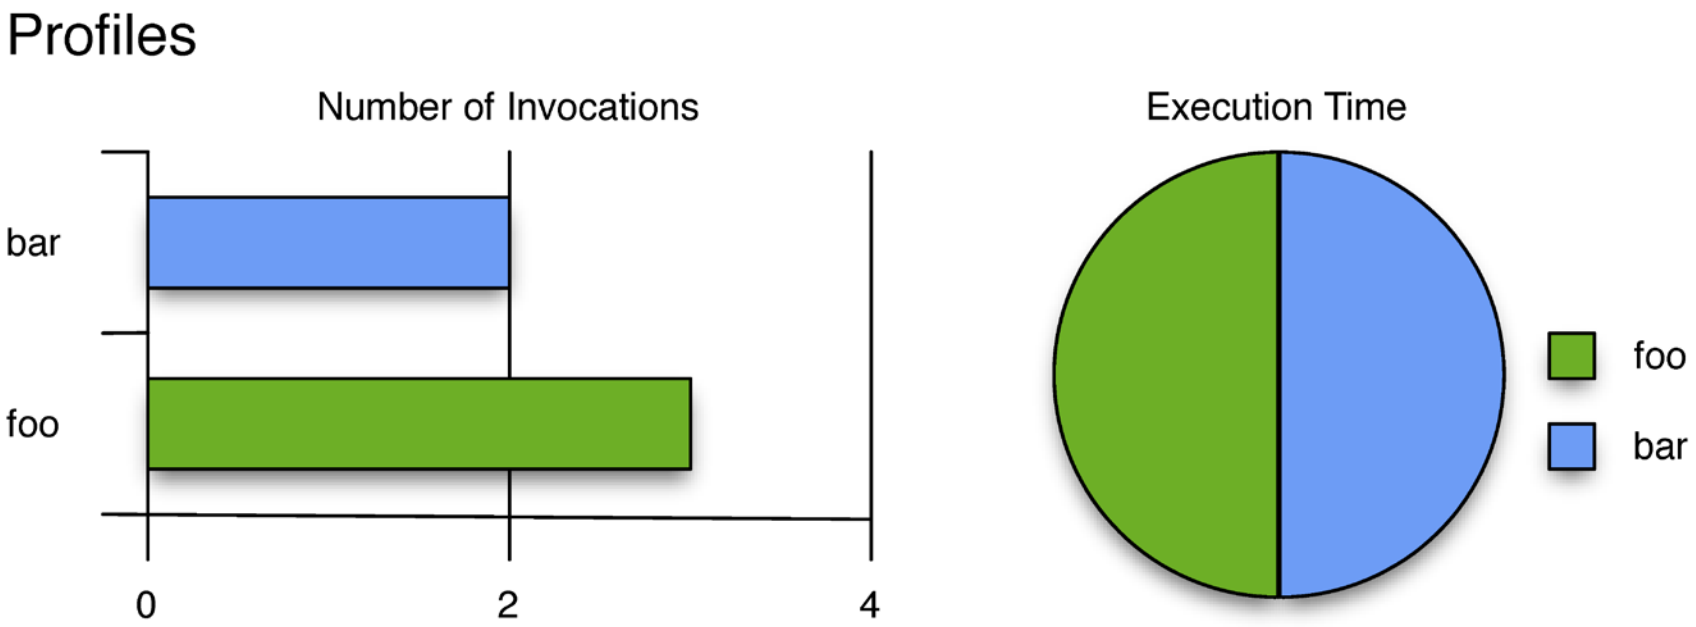

How to get it?

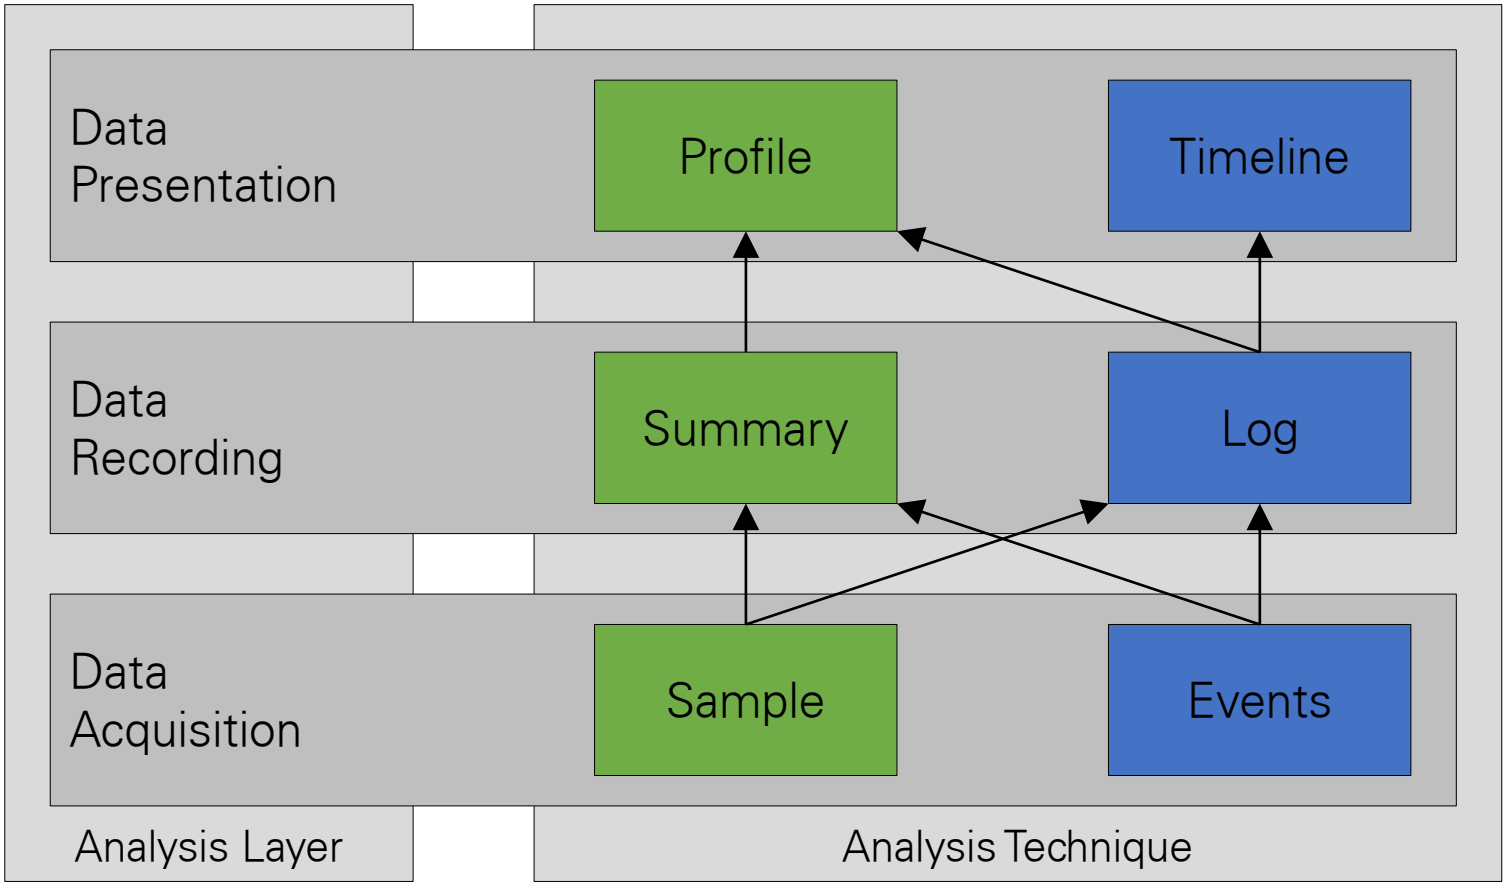
Data Presentation
Profile
Timeline
Data Recording
Summary
Log
Data Acquisition
Sample
Events
Analysis Layer
Analysis Technique

Sampling vs. Tracing

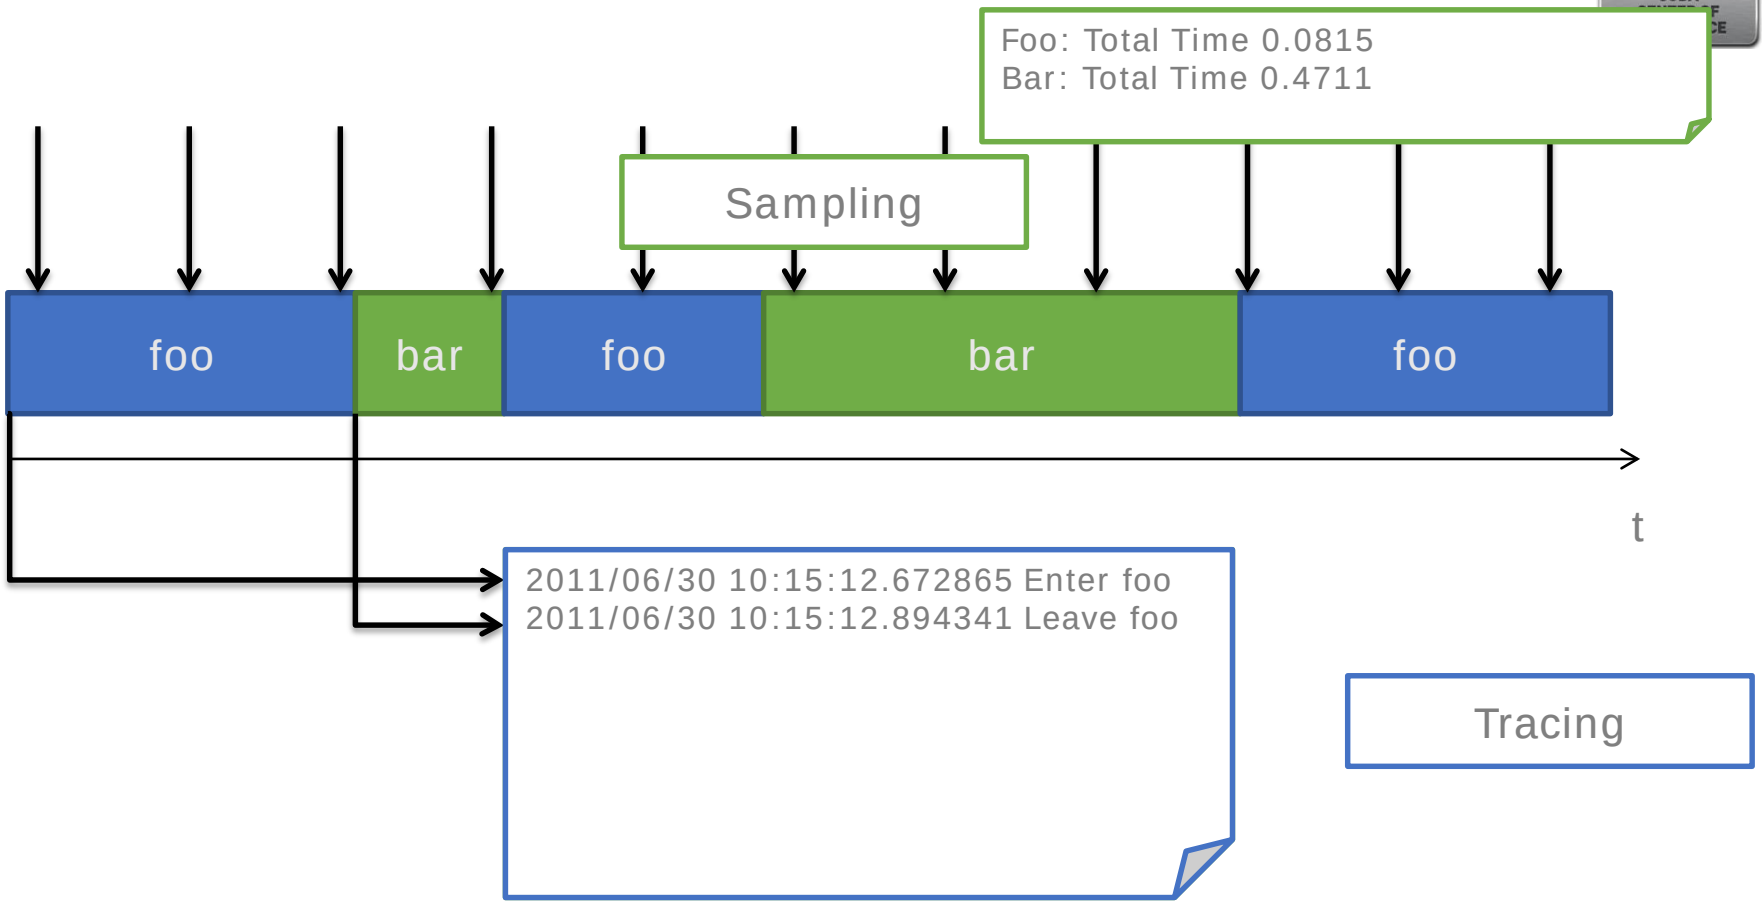

Using Real Tools on Real Applications

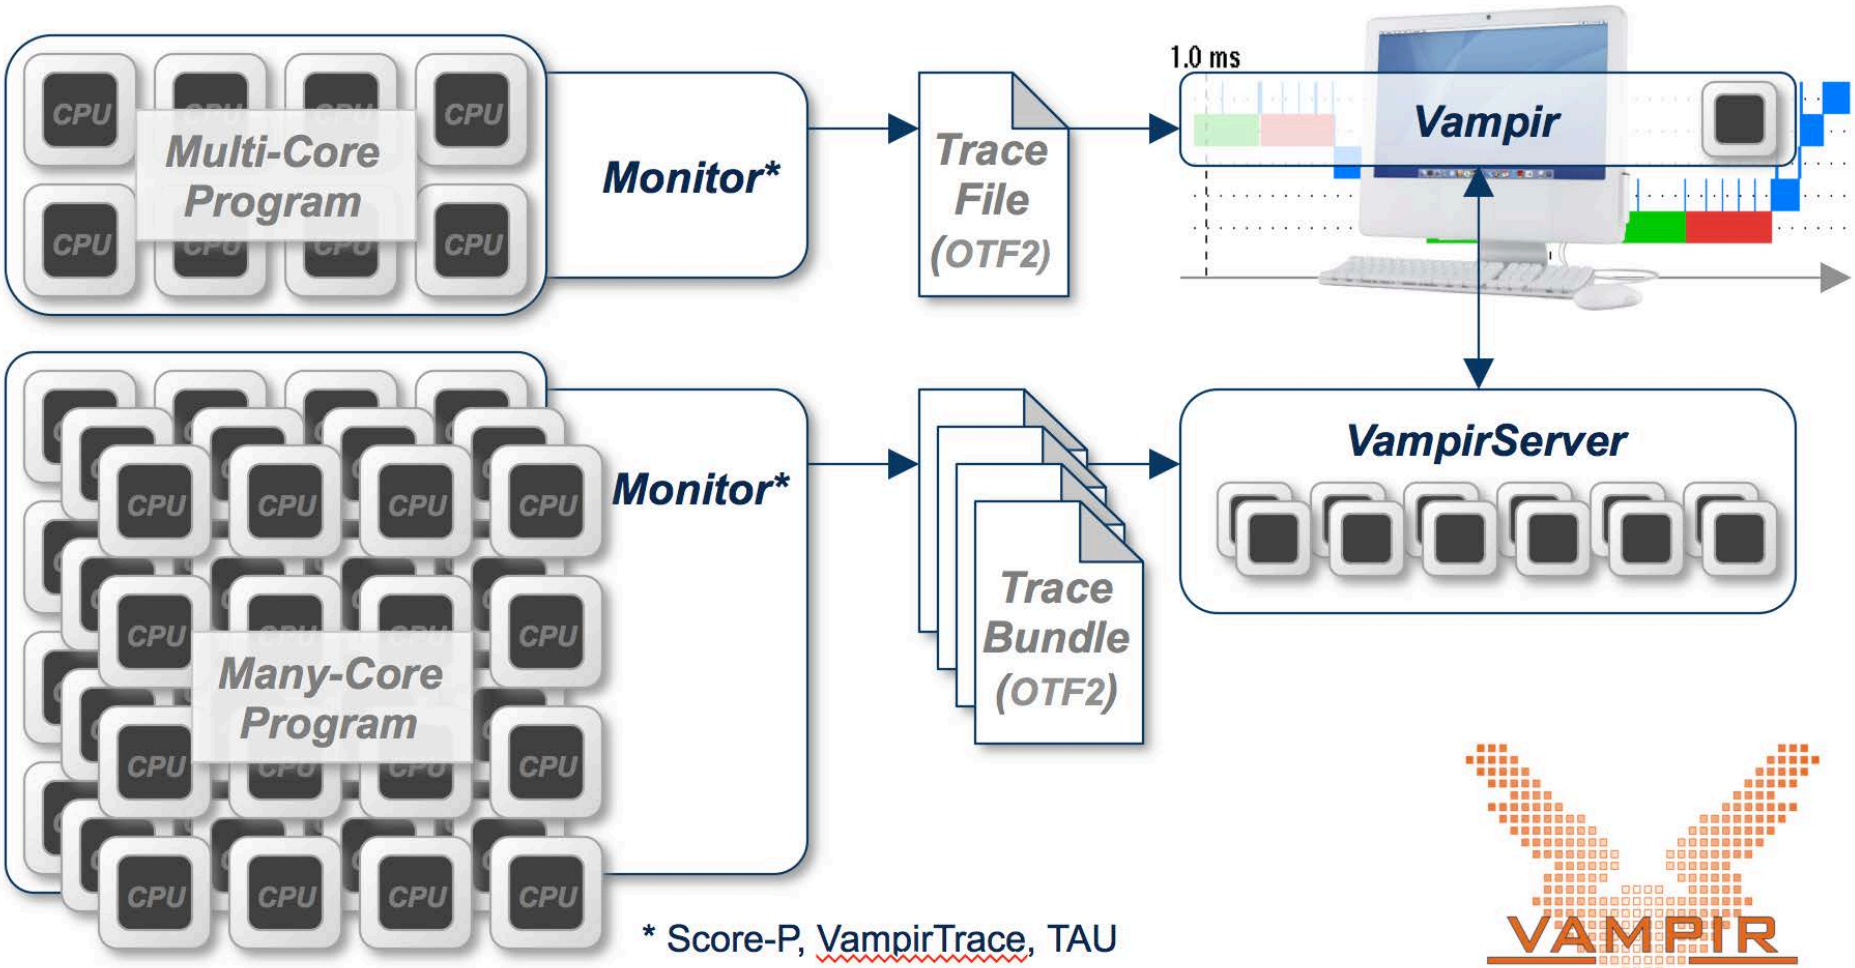

Score-P as a Collaboration Between Tool Providers

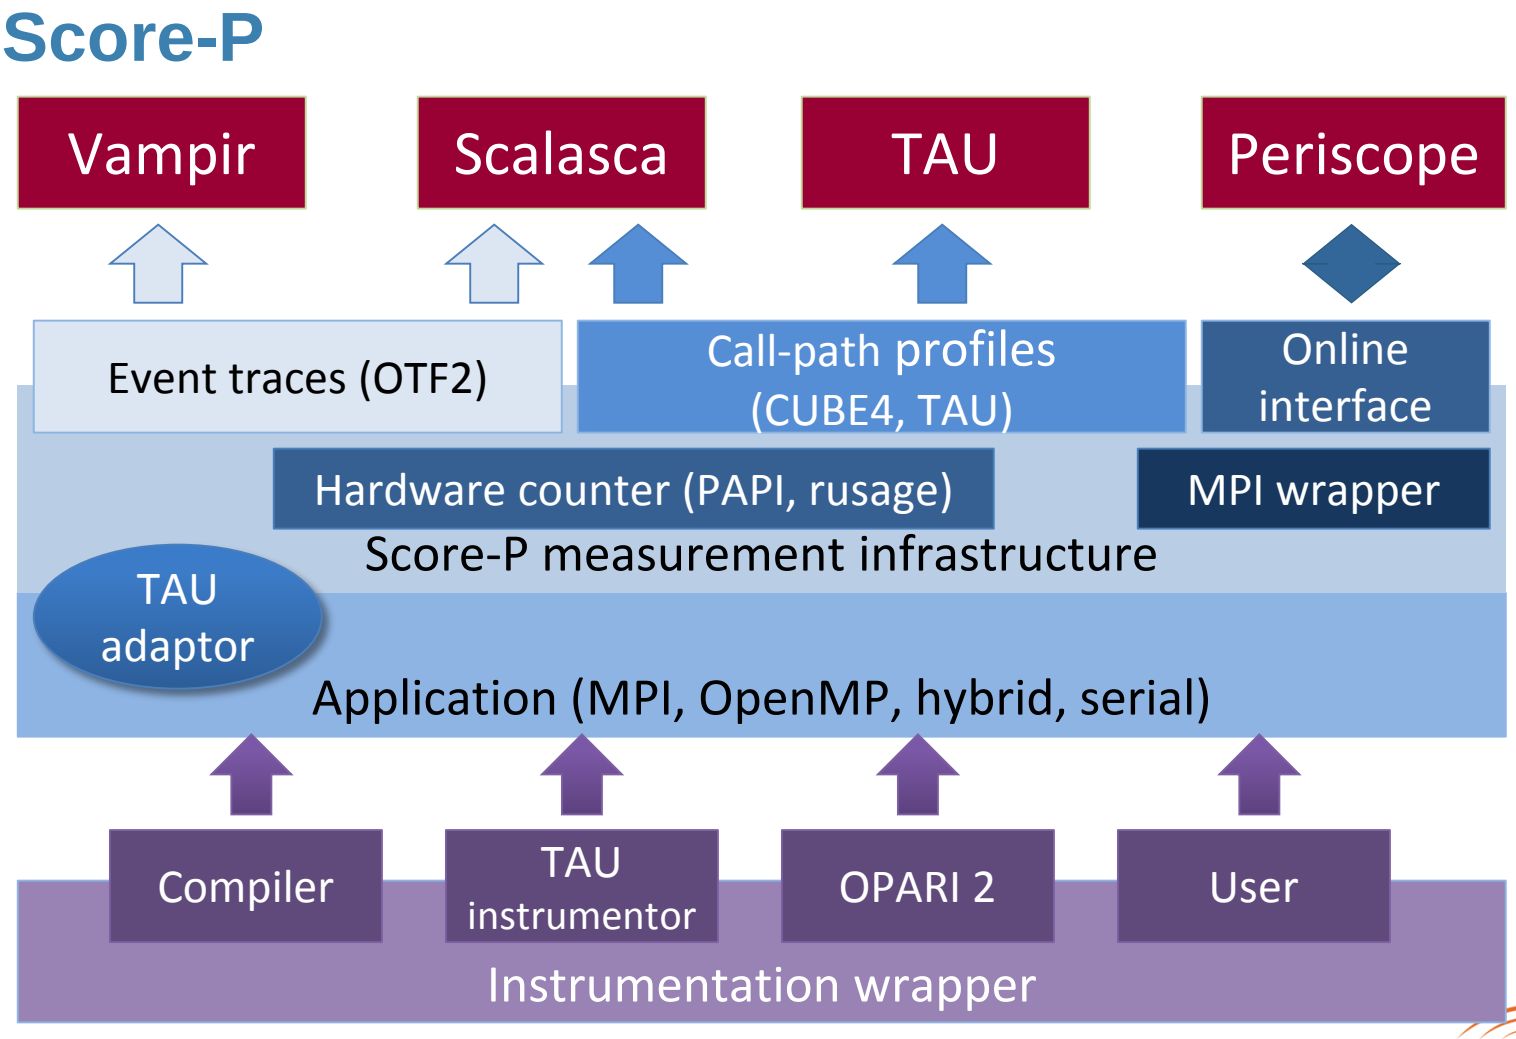

Vampir 8 as an Example for Performance Data Visualization

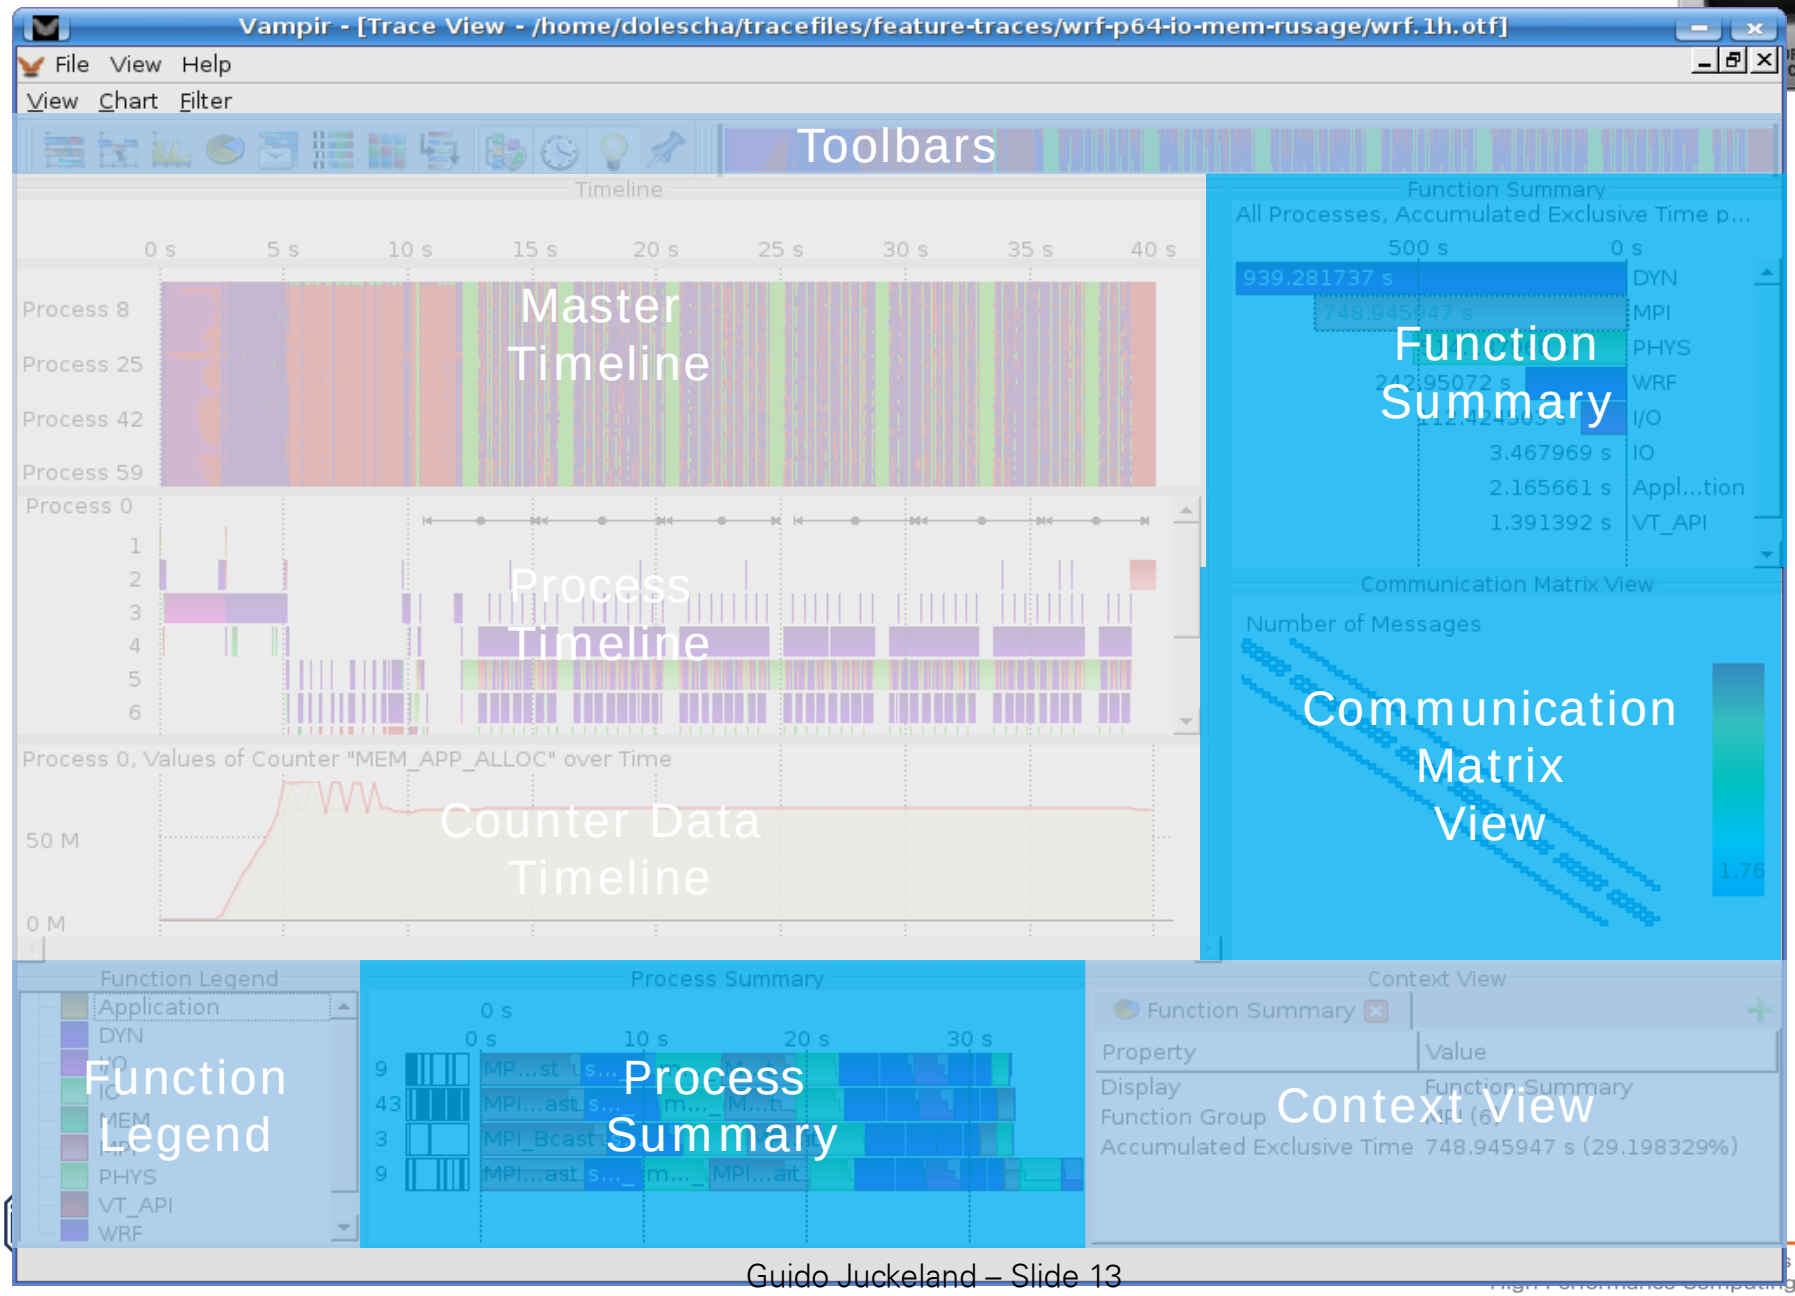

USING PERFORMANCE TOOLS FOR ACCELERATORS

The Accelerator Challenge: Asynchronity

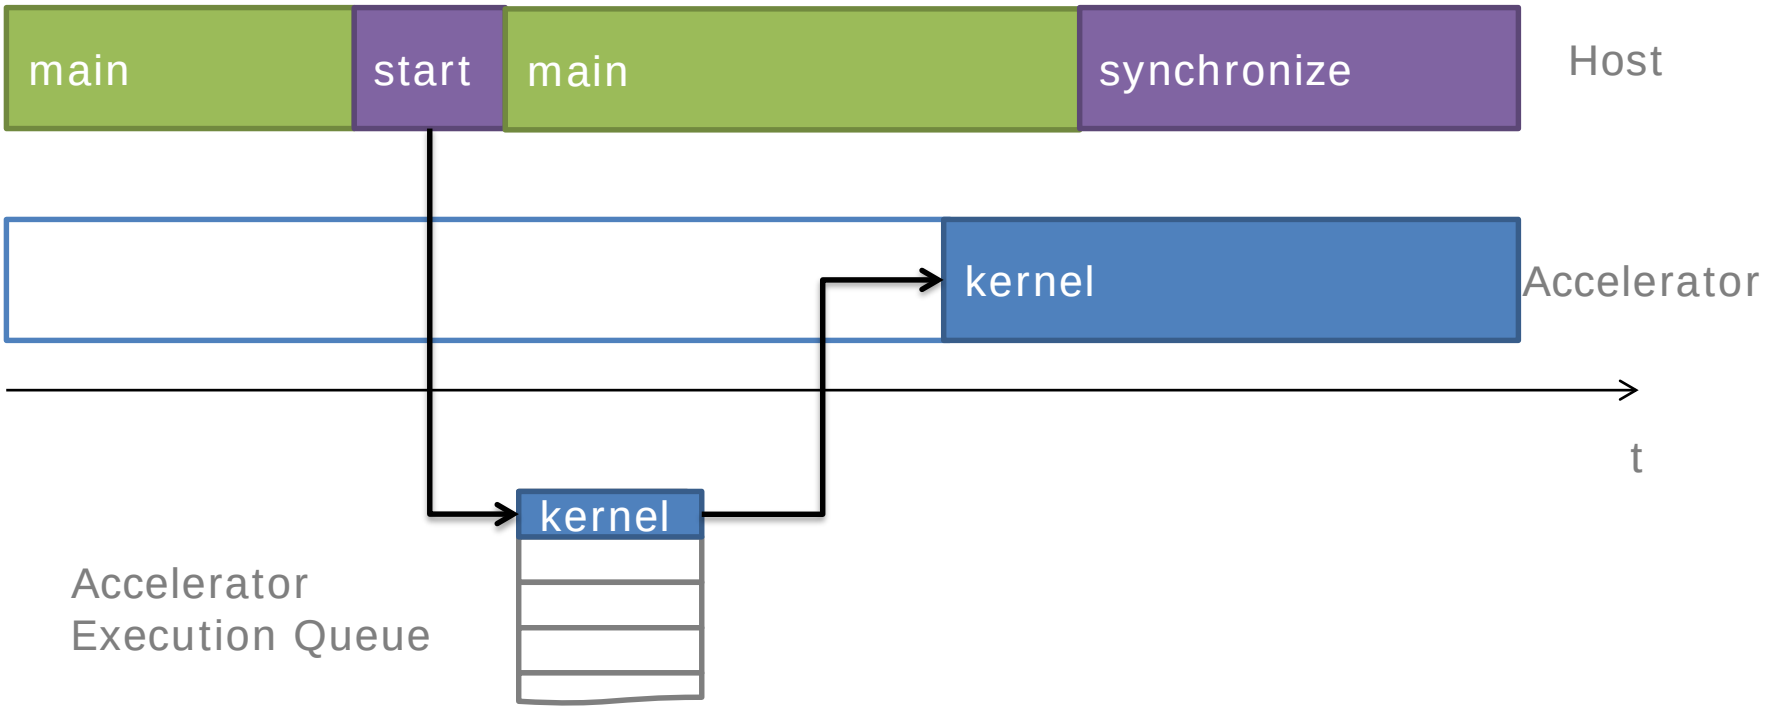

Working Together with the Vendor: CUPTI

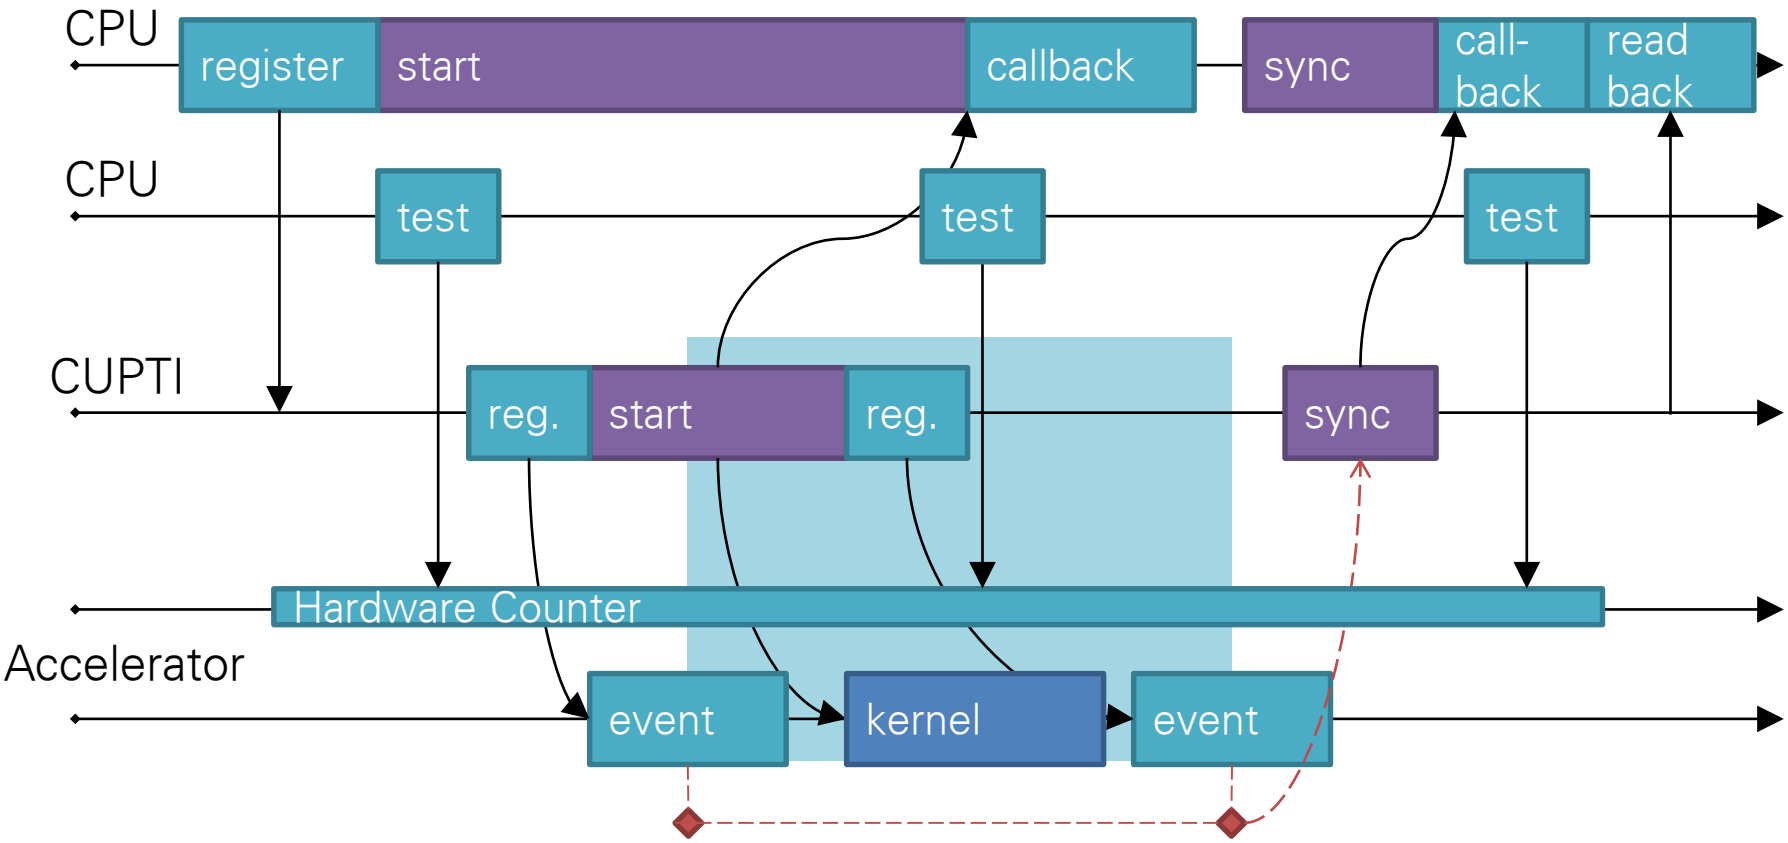

Similar things possible for OpenCL

What About Directive Based Approaches?

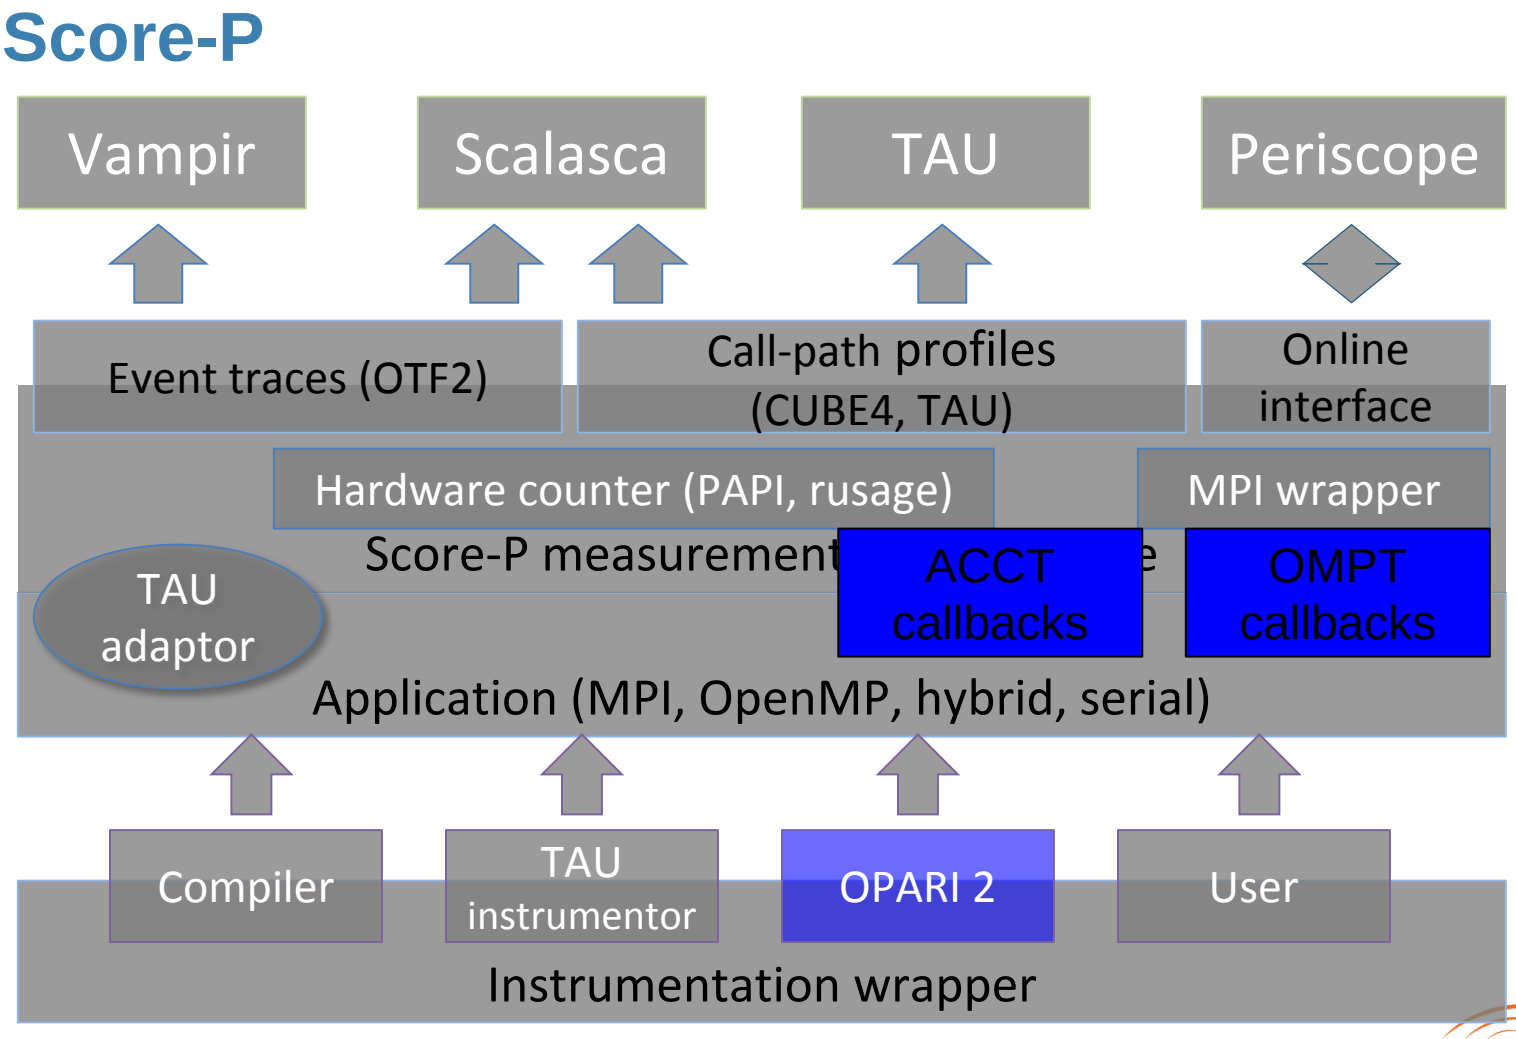

Comparing Monitoring Tool Capabilities

	MPI	Threads	Accel-erator	Monitoring Method	Scala-bility
Vendor Tools				Event + Sample Summary and Log	
VampirTrace / Score-P				Event + Sample Log	
TAU				Ereignis + Sample Aufzeichnung	
HPCtoolkit				Event + Sample Summary and Log	
IPM				Event Summary	
CEPBA MPItrace				Event Log	
PAPI				Sample Summary	
GPU Ocelot				Event Summary	

Looking at Overhead: PIConGPU using 512 GPUs

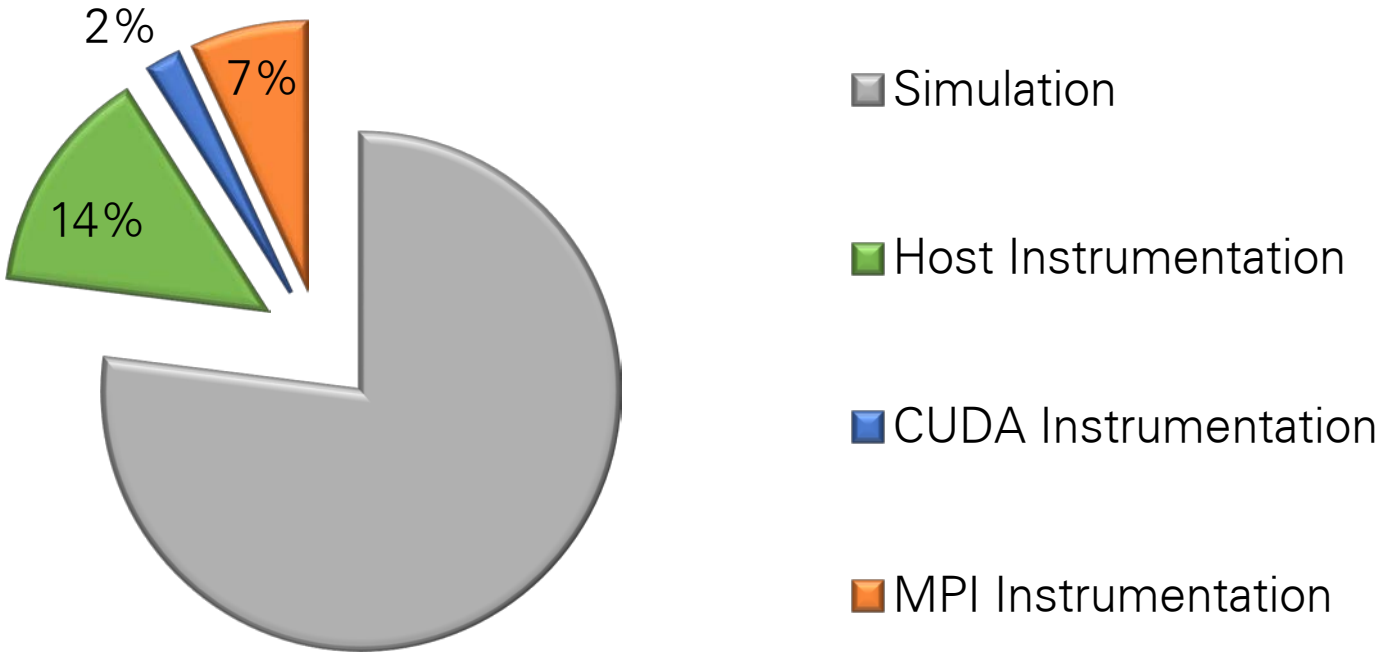

EXAMPLES

Looking at Multi-hybrid Application

Single GPU Implementation (5 years ago)

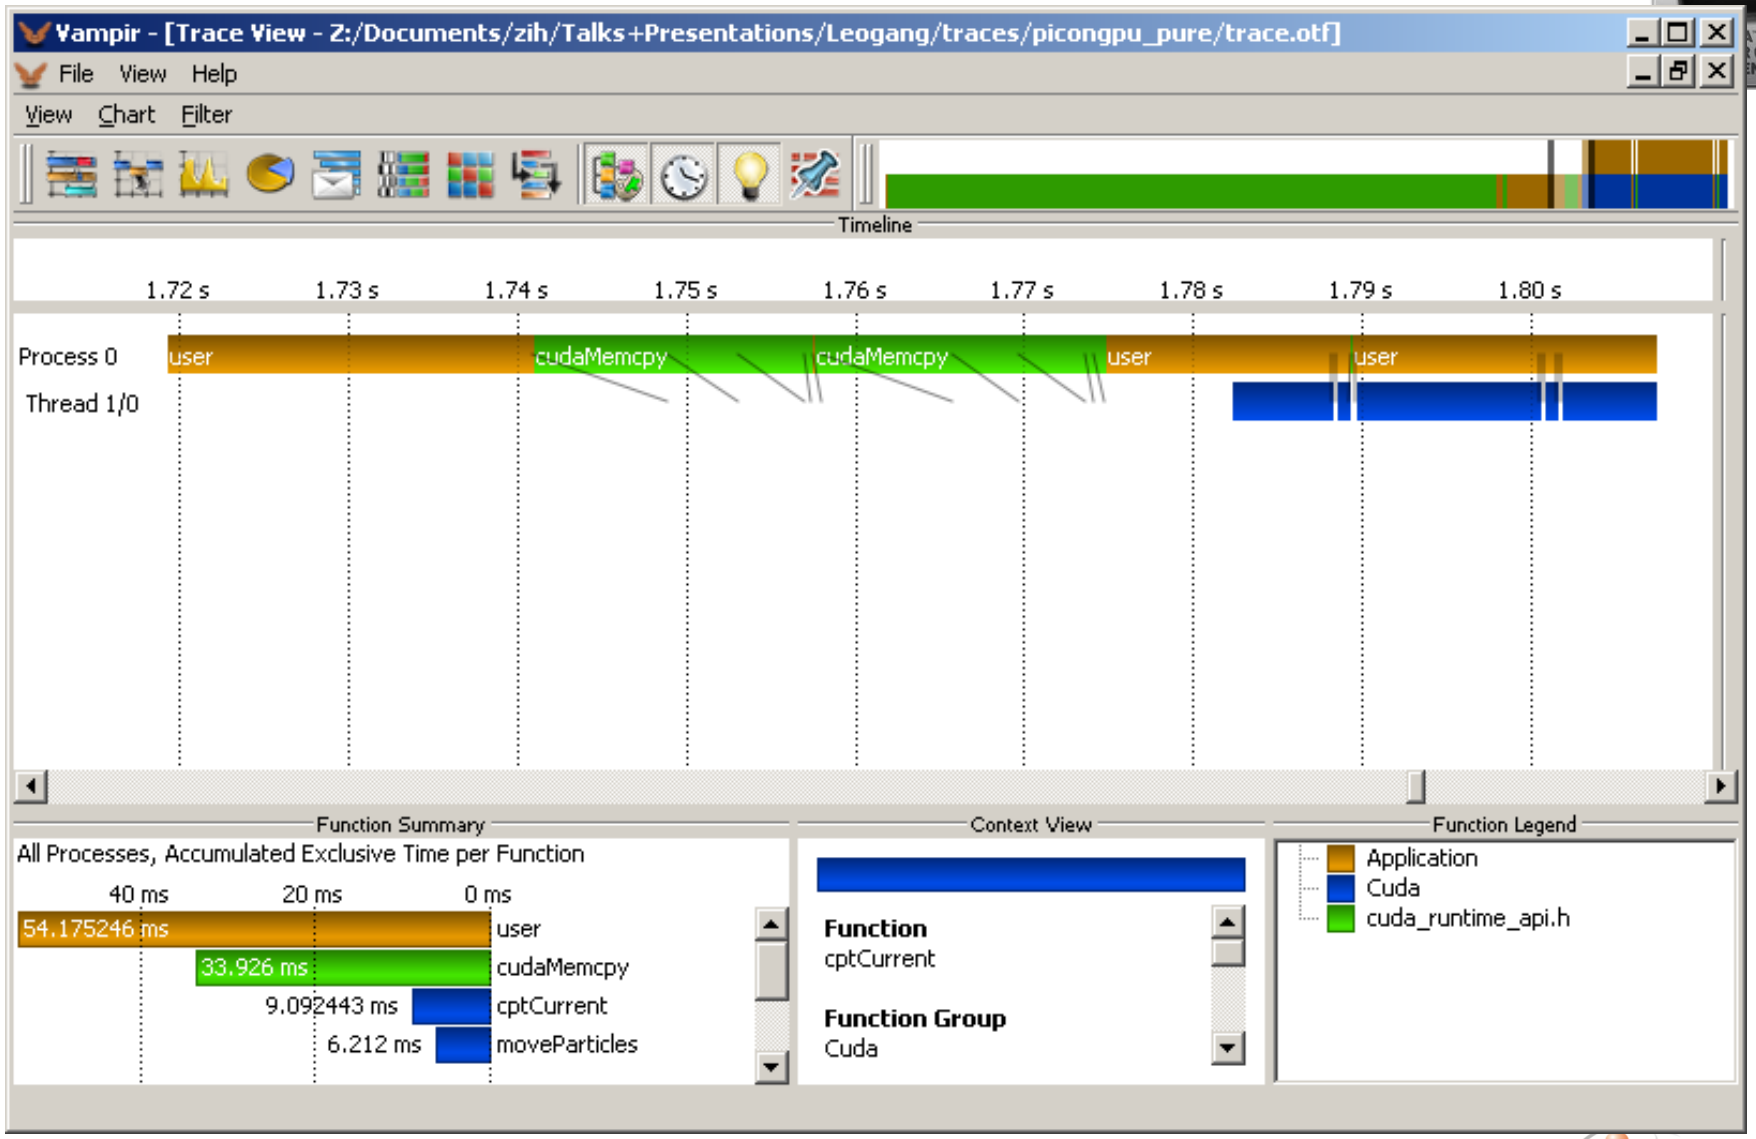

Inter-GPU-Communication with synchronous MPI

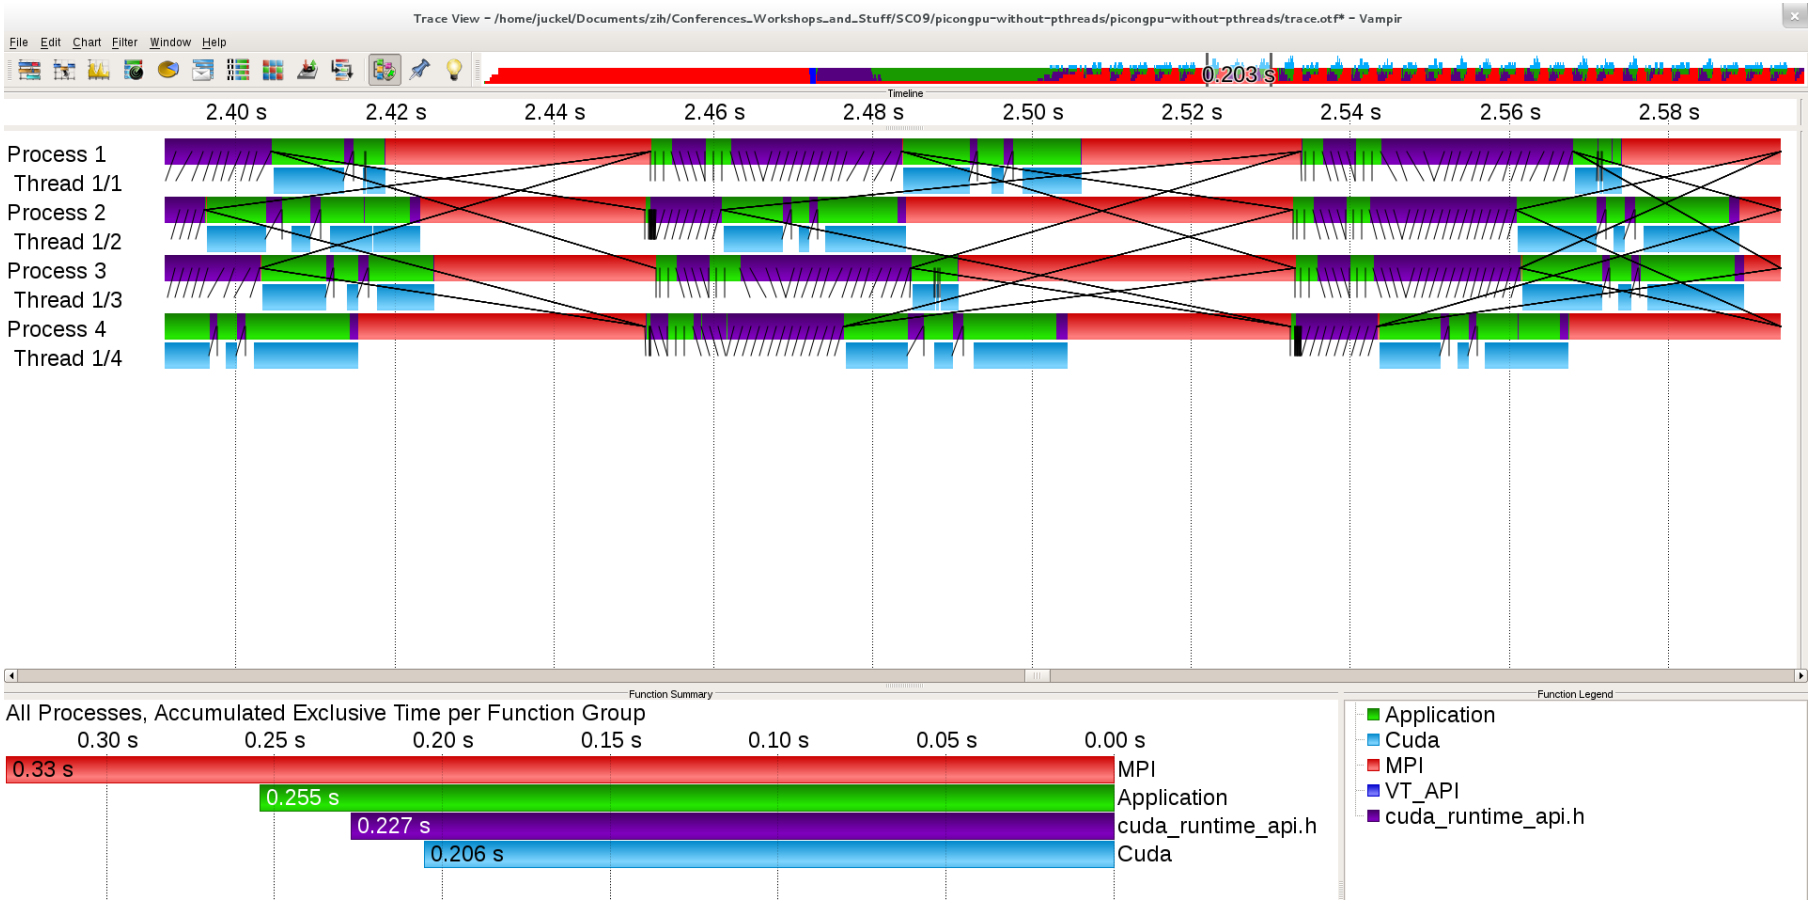

Impact of vectorized Kernels and asynchronous Communication

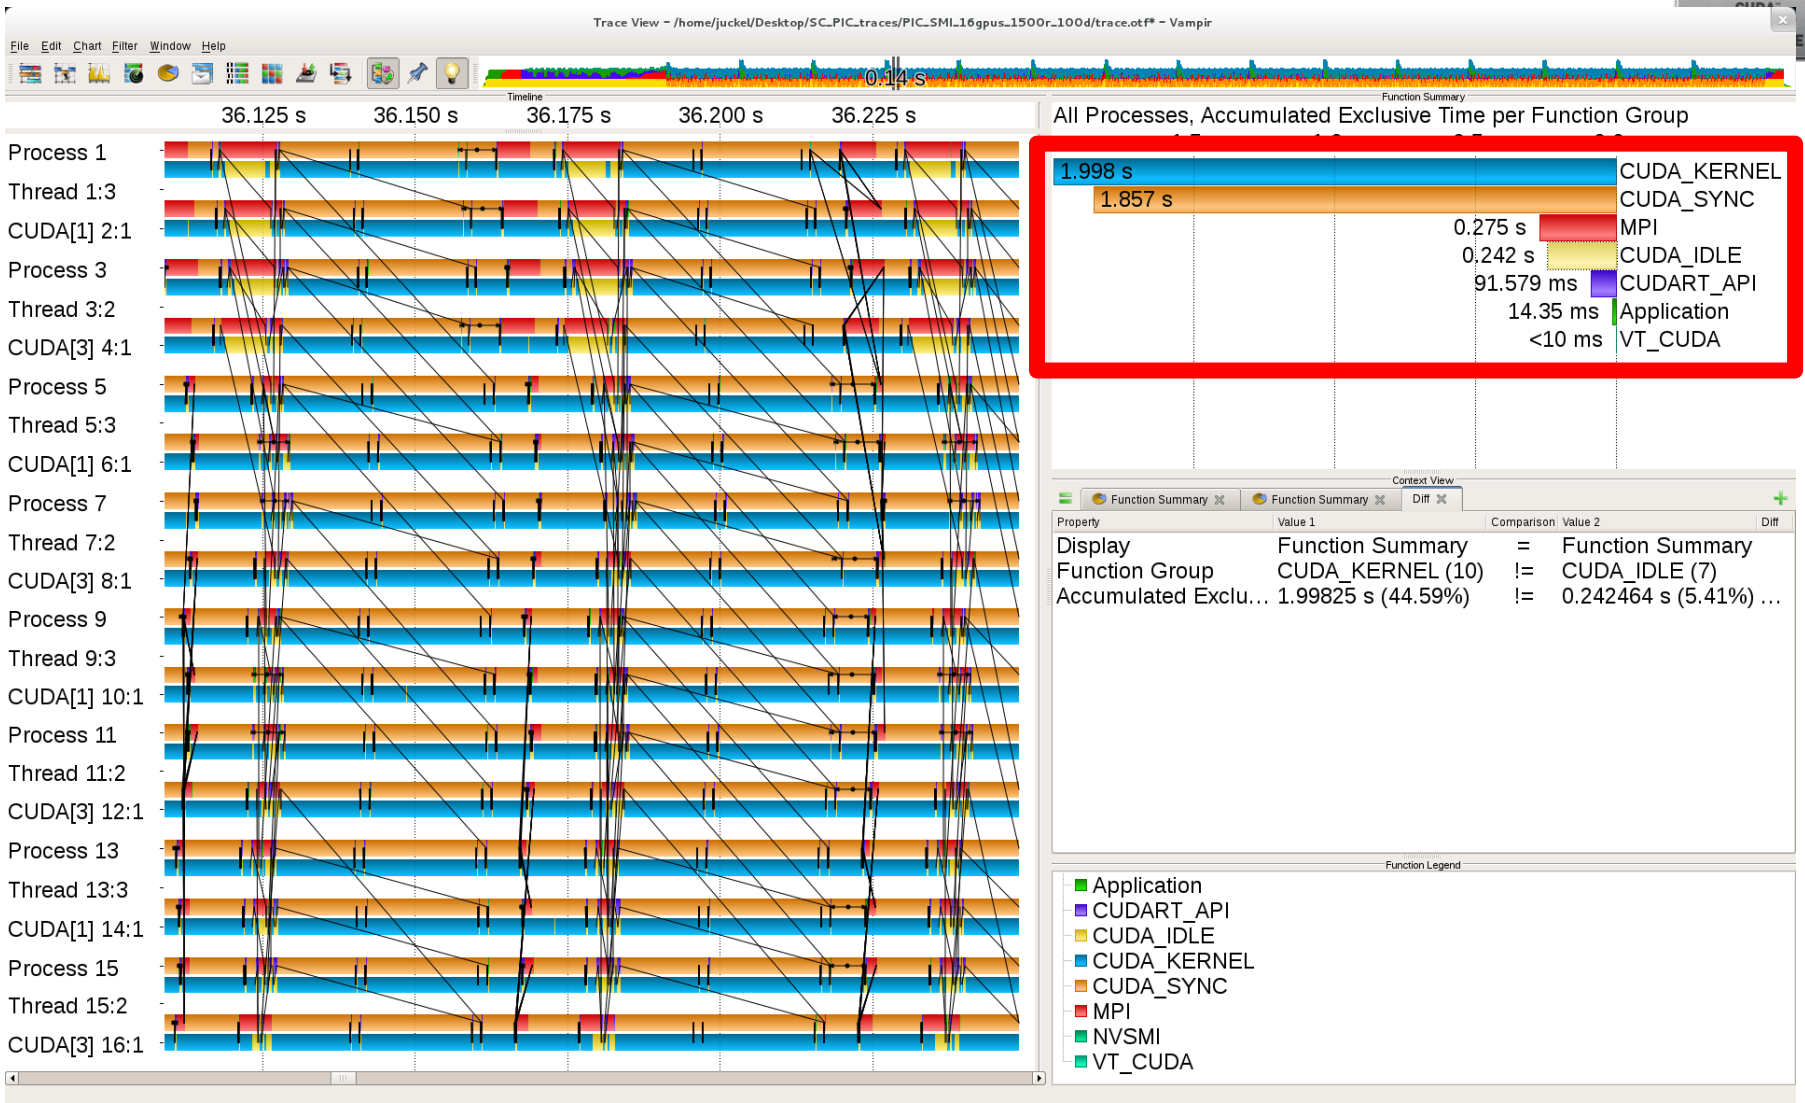

Concurrent Kernel Execution and Communication

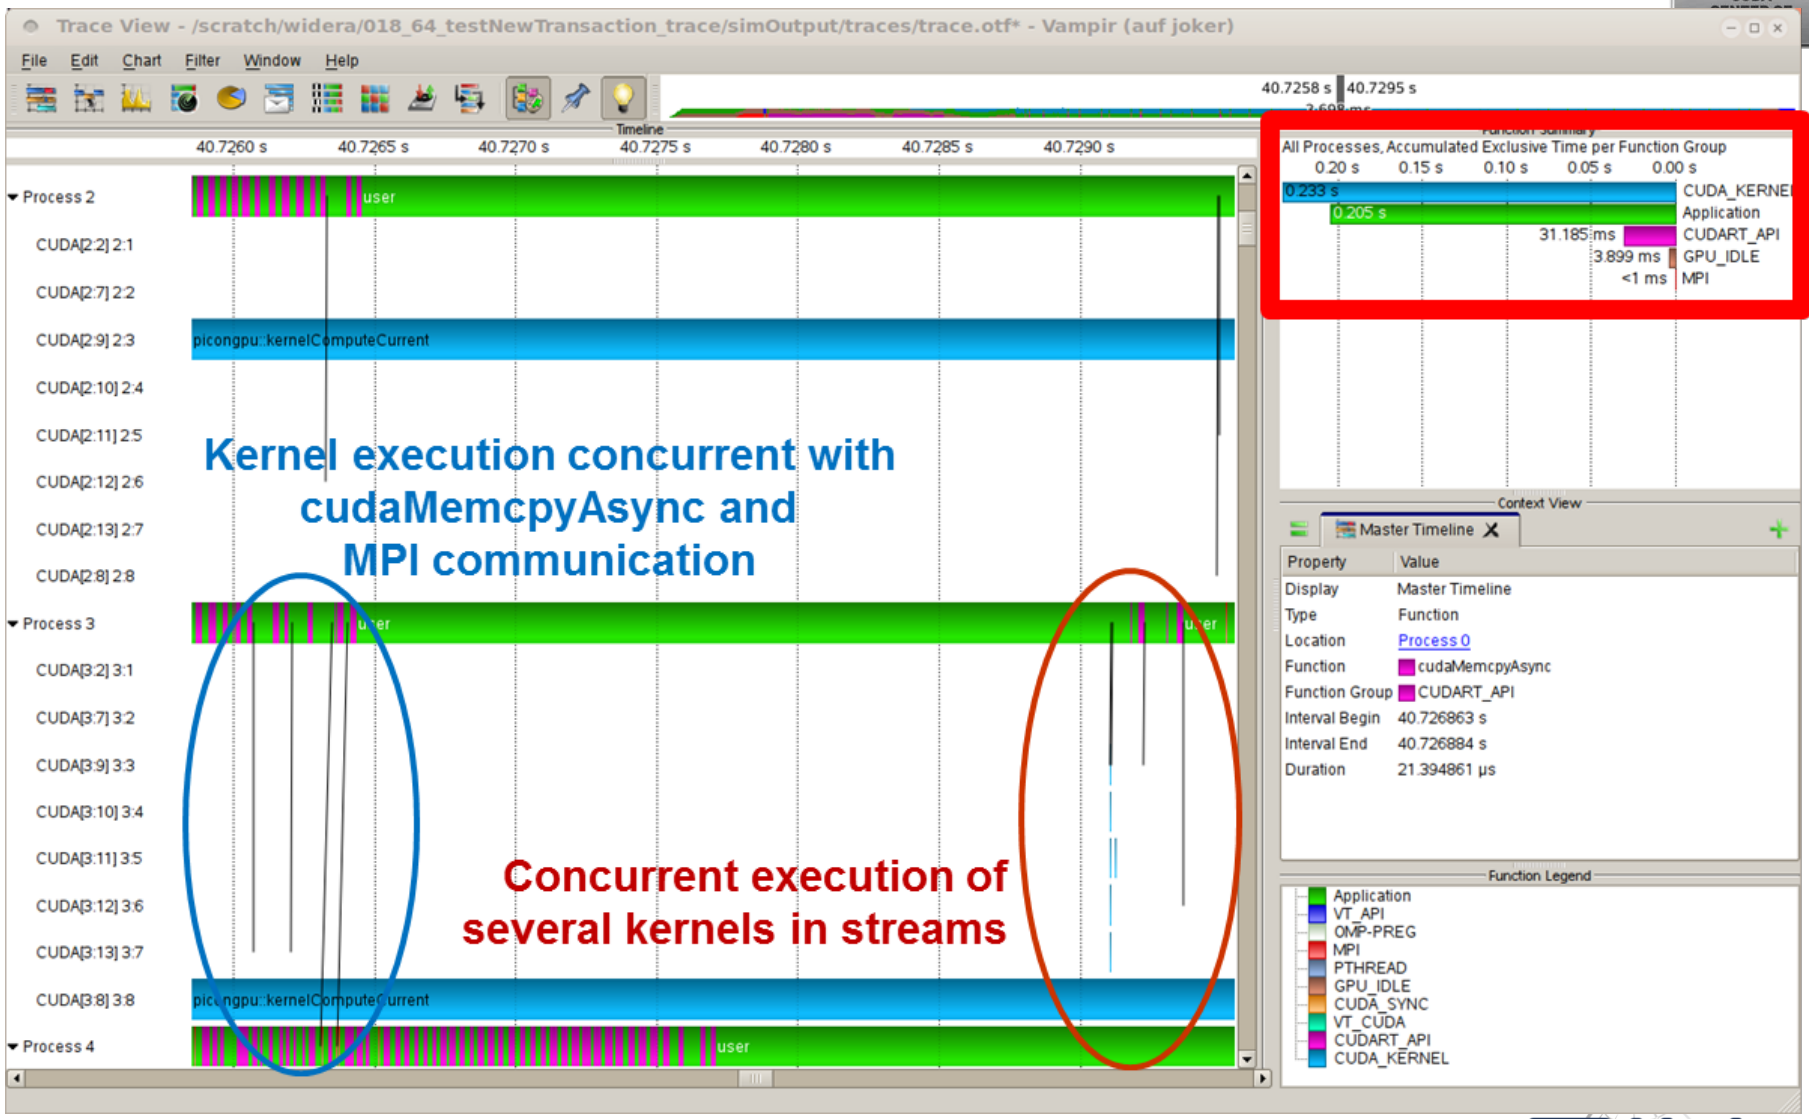

Going to Large GPU Counts

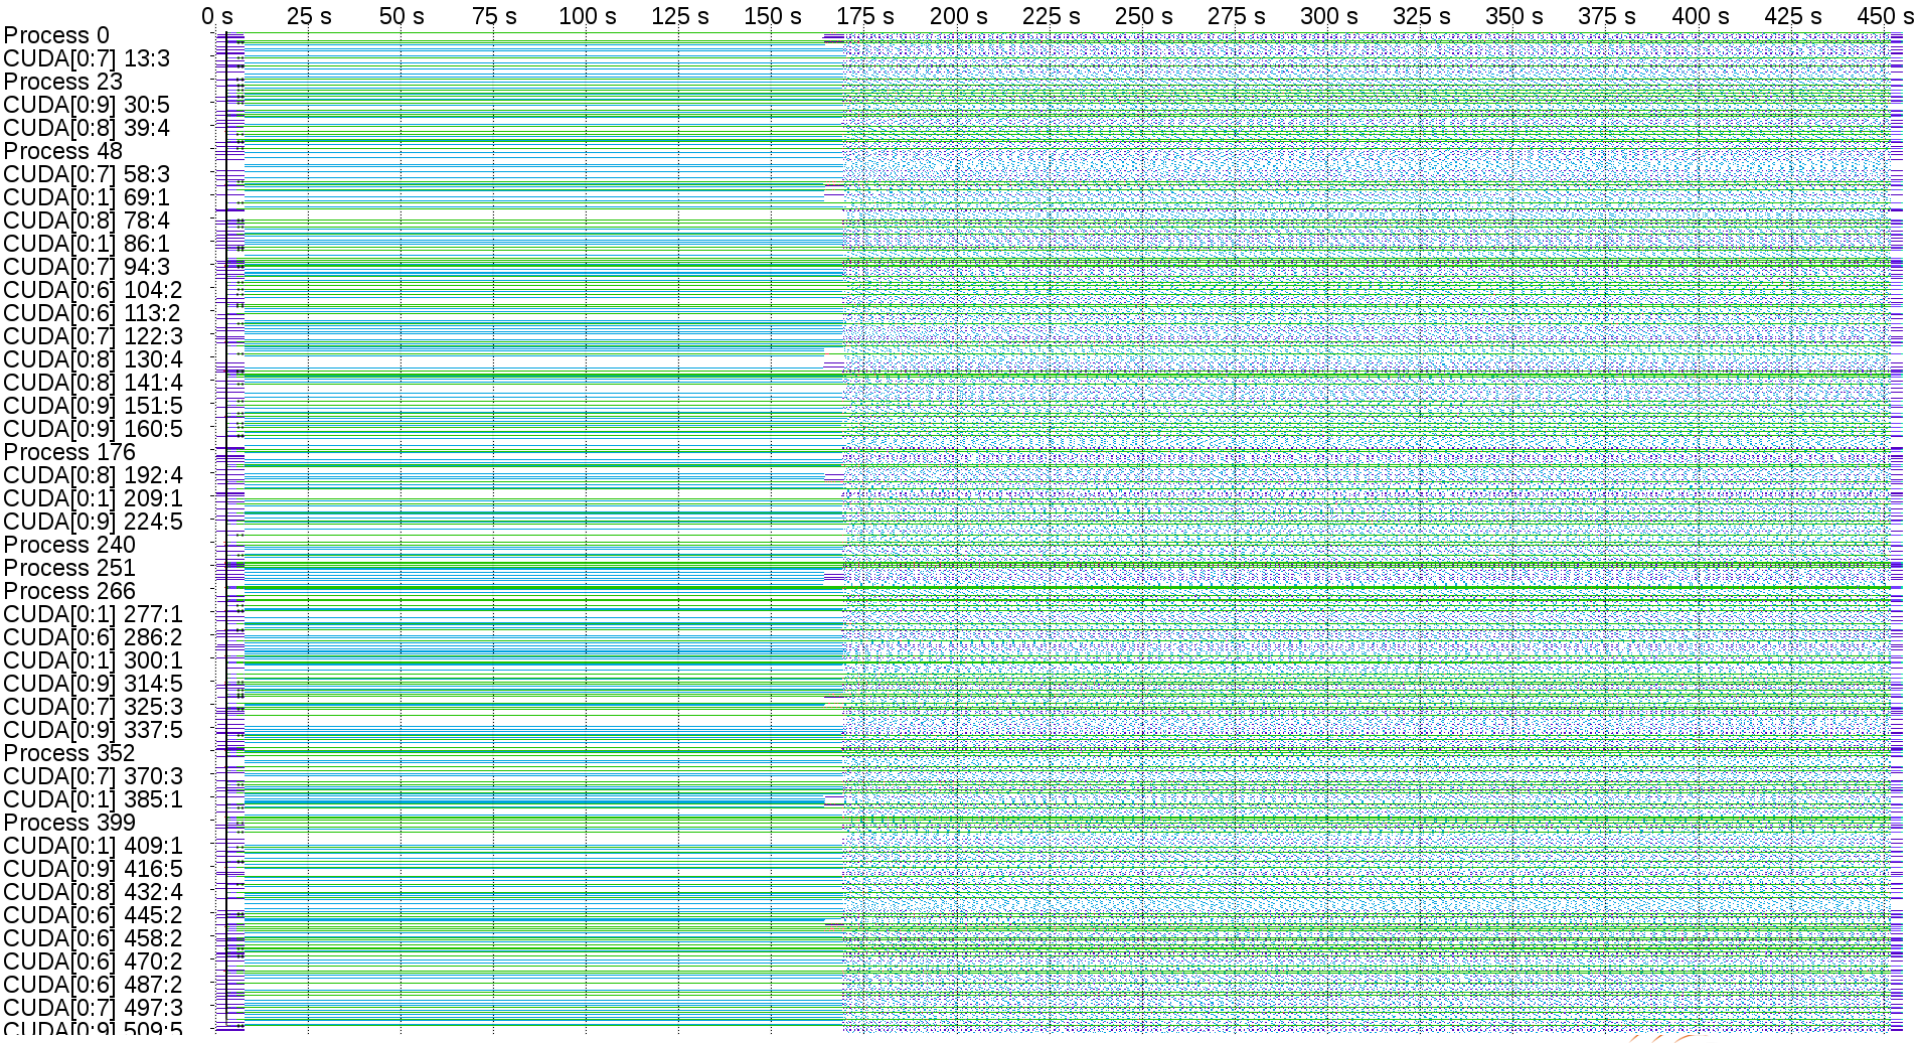

Filtering

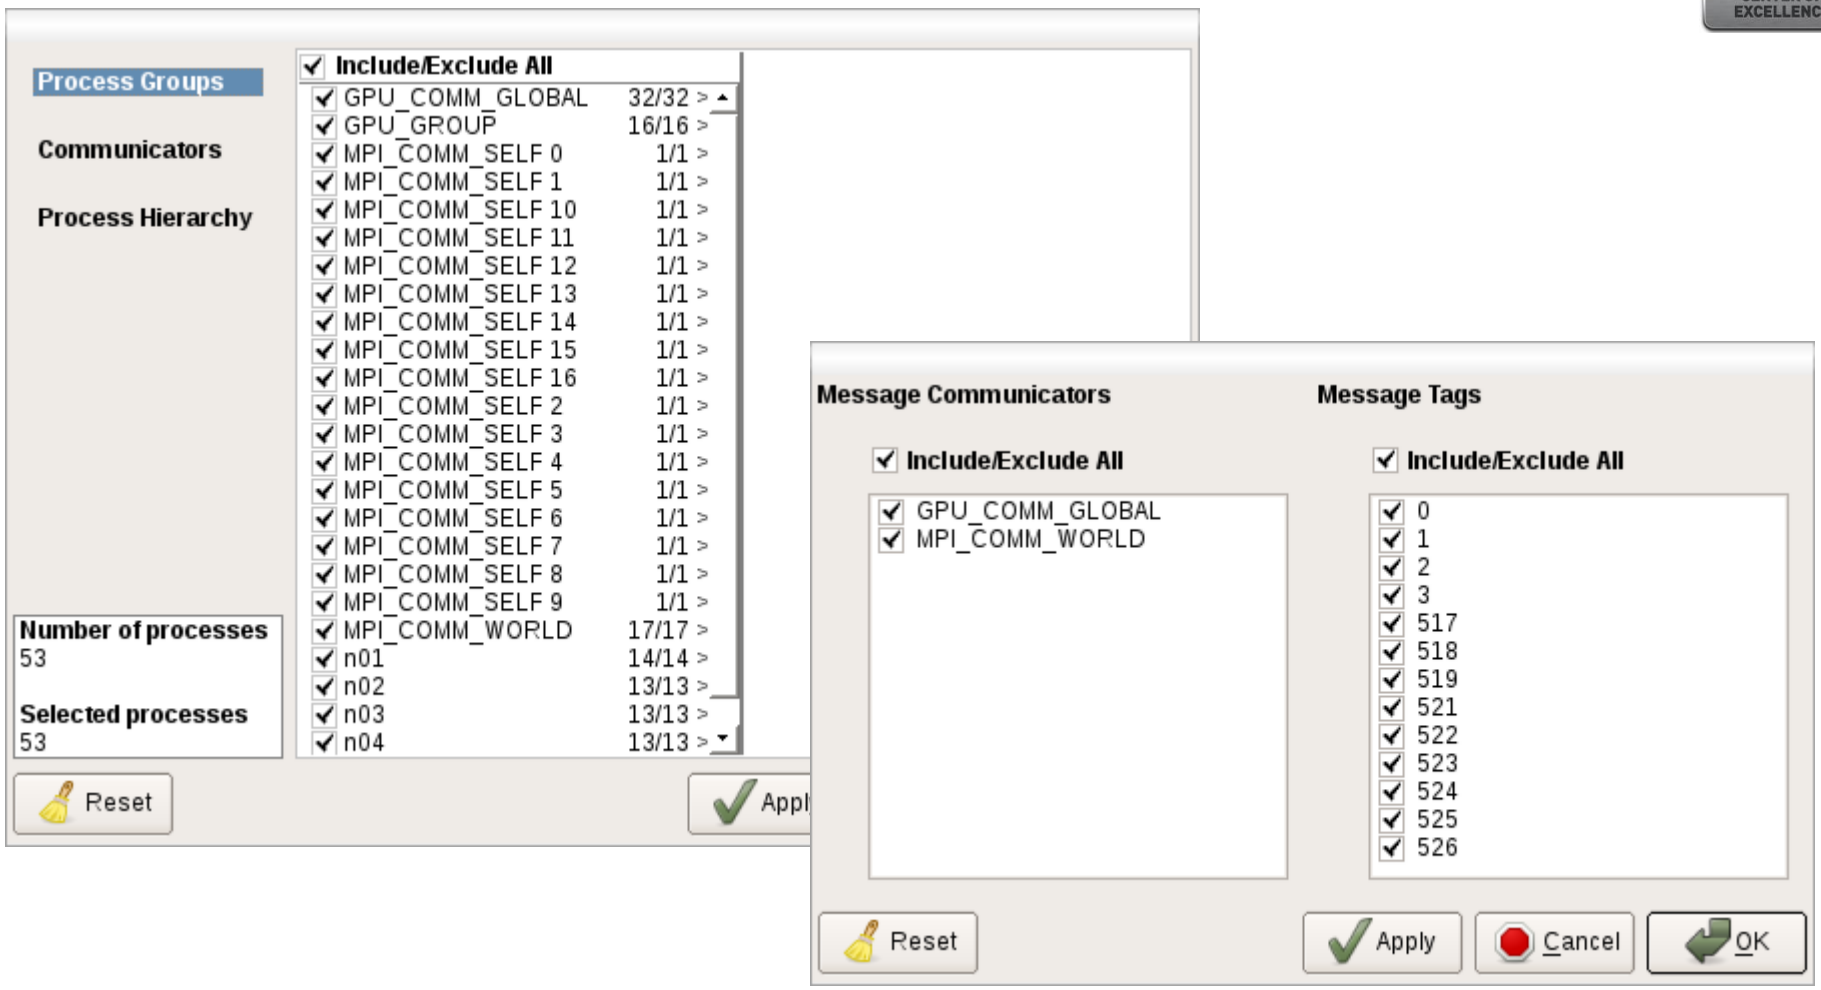
Process Groups
Communicators
Process Hierarchy
Include/Exclude All
GPU_COMM_GLOBAL 32/32 >
GPU_GROUP 16/16 >
MPI_COMM_SELF 0 1/1 >
MPI_COMM_SELF 1 1/1 >
MPI_COMM_SELF 10 1/1 >
MPI_COMM_SELF 11 1/1 >
MPI_COMM_SELF 12 1/1 >
MPI_COMM_SELF 13 1/1 >
MPI_COMM_SELF 14 1/1 >
MPI_COMM_SELF 15 1/1 >
MPI_COMM_SELF 16 1/1 >
MPI_COMM_SELF 2 1/1 >
MPI_COMM_SELF 3 1/1 >
MPI_COMM_SELF 4 1/1 >
MPI_COMM_SELF 5 1/1 >
MPI_COMM_SELF 6 1/1 >
MPI_COMM_SELF 7 1/1 >
MPI_COMM_SELF 8 1/1 >
MPI_COMM_SELF 9 1/1 >
MPI_COMM_WORLD 17/17 >
n01 14/14 >
n02 13/13 >
n03 13/13 >
n04 13/13 >
Number of processes
53
Selected processes
53
Reset
Appl
Message Communicators
Include/Exclude All
GPU_COMM_GLOBAL
MPI_COMM_WORLD
Message Tags
Include/Exclude All
0
1
2
3
517
518
519
521
522
523
524
525
526
Reset
Apply
Cancel
OK

Only GPU activity

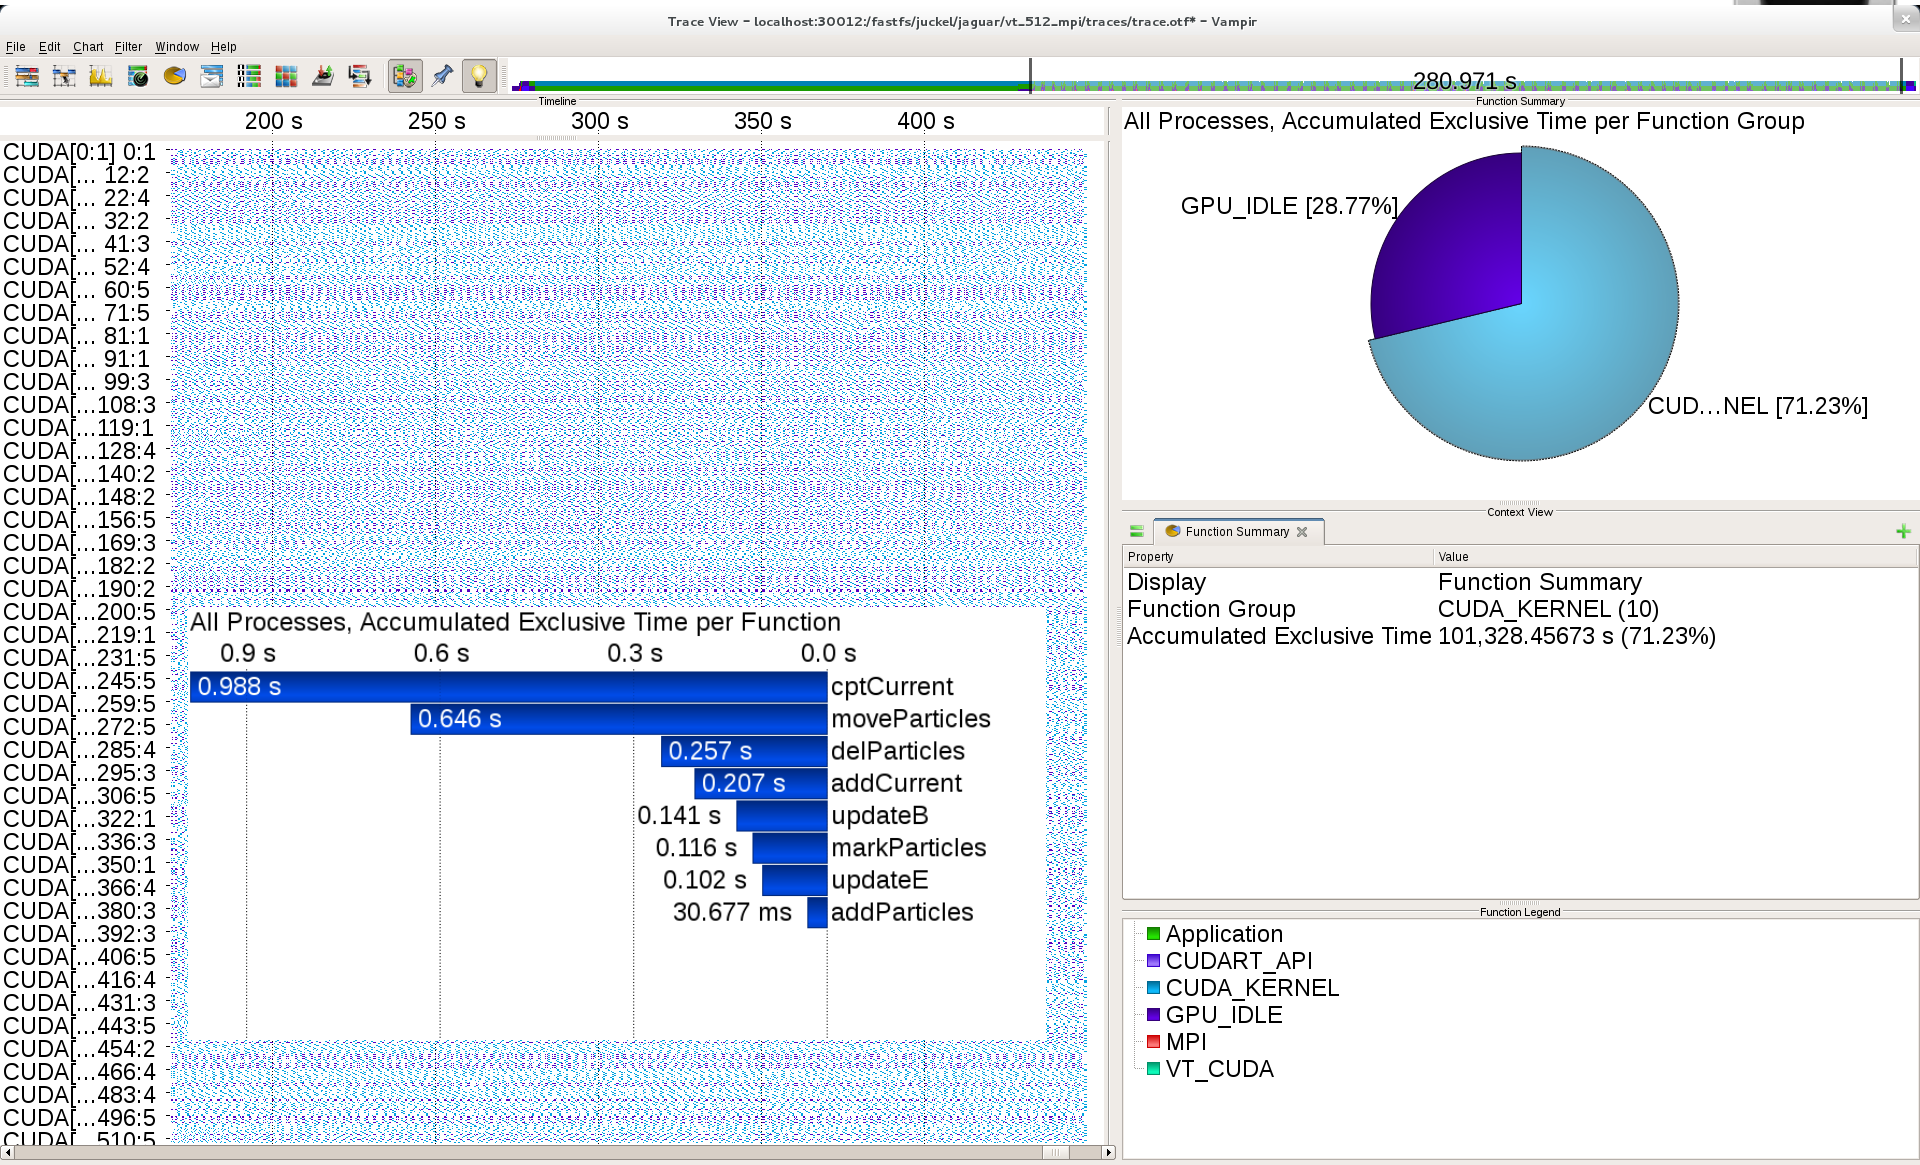

Now What About Directives?

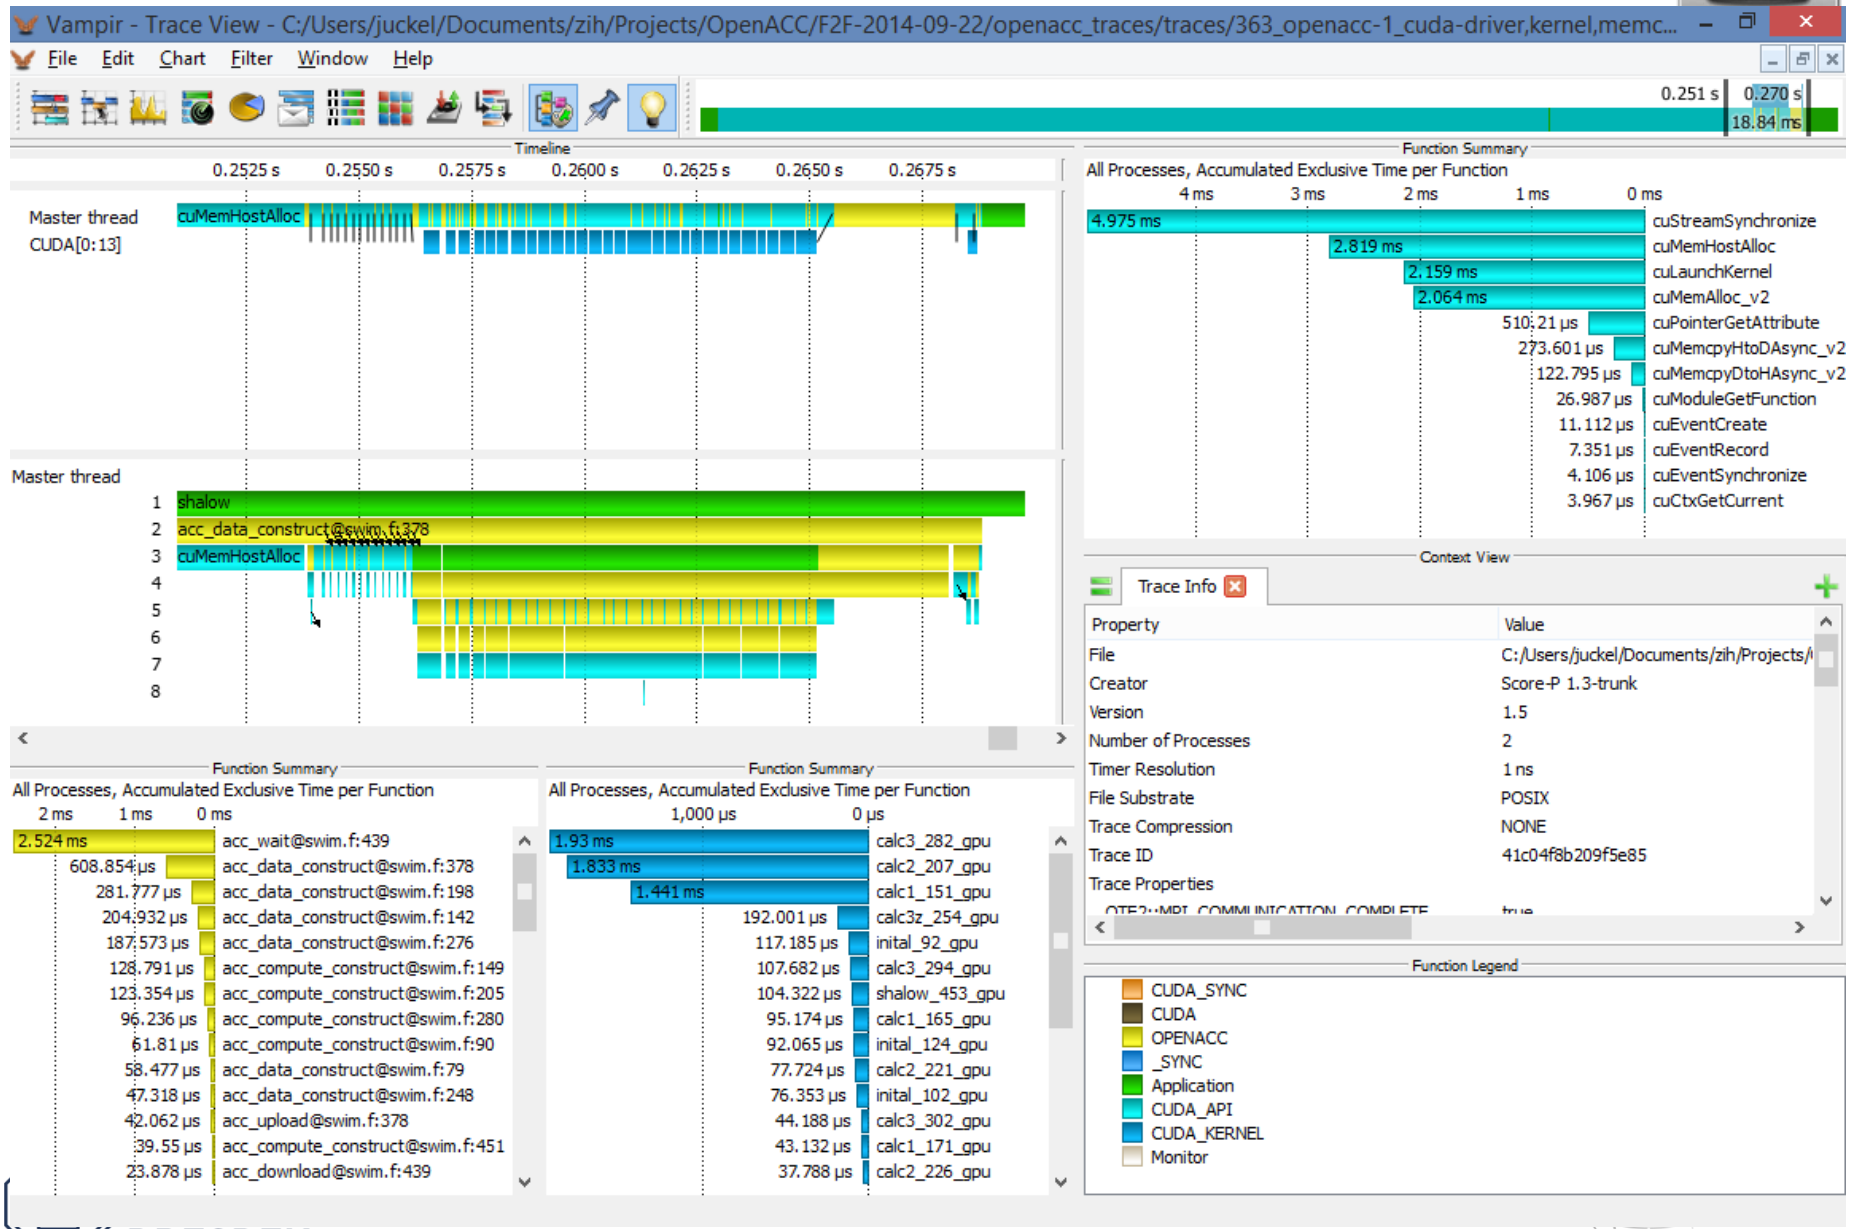

SUMMARY & OUTLOOK

Summary

All levels of parallelism visible

- Inter-node (MPI, SHMEM)
- Intra-node (OpenMP, pthreads)
- Accelerators (CUDA, OpenCL)

Multiple highly scalable analysis tools available

- Scalasca
- Vampir

Experts available on-site

- You are running out of excuses… ;-)

Outlook

Comittee work

- OpenACC (profiler interface)
- OpenMP (OMPT)
- Score-P group (finding usable solutions)

Critical Path Analysis

- Blaming the right application parts

Questions

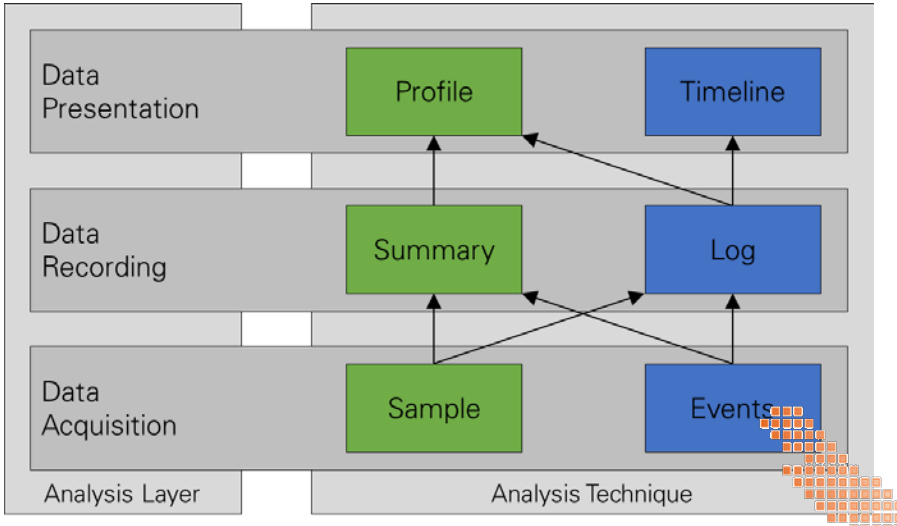
Data Presentation
Profile
Timeline
Data Recording
Summary
Log
Data Acquisition
Sample
Events
Analysis Layer
Analysis Technique

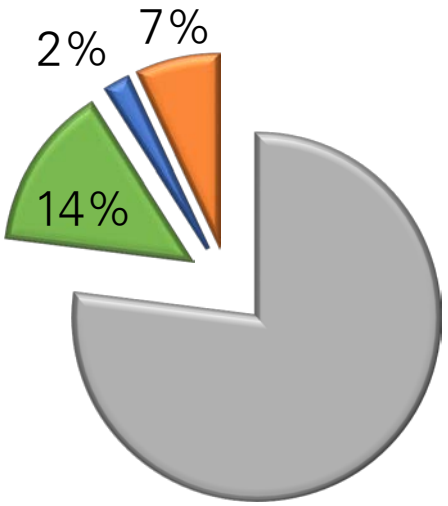
2%
7%
14%

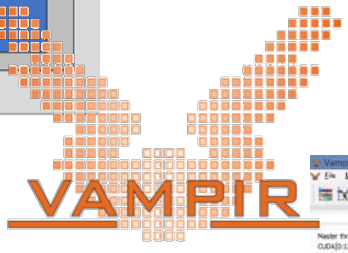
VAMPIR

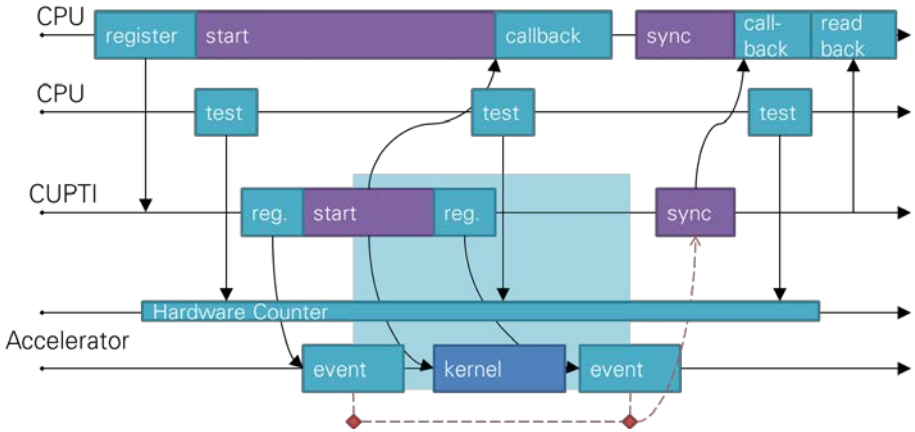
CPU
register
start
callback
sync
call-back
read back
CPU
test
test
test
CUPTI
reg.
start
reg.
sync
Hardware Counter
Accelerator
event
kernel
event